PNAS
www.pnas.org

Supplementary Information for

How a well-adapting immune system remembers

Andreas Mayer, Vijay Balasubramanian, Aleksandra M. Walczak, Thierry Mora

Correspondence: Aleksandra M. Walczak awalczak@lpt.ens.fr, Thierry Mora tmora@lps.ens.fr

This PDF file includes:

Supporting Information Text

1. The immigration-drift model of pathogen dynamics

A. Steady-state distribution and the backward equation. The steady-state distribution of the pathogen dynamics (Eq. 11) is a Dirichlet distribution with the parameter vector $\boldsymbol{\theta}$ (1),

$$\rho_s(\boldsymbol{Q}) = \frac{1}{Z(\boldsymbol{\theta})} \prod_i Q_i^{\theta_i - 1} =: \mathcal{D}(\boldsymbol{Q}, \boldsymbol{\theta}). \quad [1]$$

The normalizing constant is the multivariate Beta function, defined in terms of Gamma functions as $Z(\boldsymbol{\theta}) = \prod_i \Gamma(\theta_i)/\Gamma(|\boldsymbol{\theta}|)$.

To obtain a solution to Eq. 11 we write $\rho(\boldsymbol{Q}, t)$ as the product of the steady-state distribution $\rho_s(\boldsymbol{Q})$ with a time-varying function $f(\boldsymbol{Q}, t)$

$$\rho(\boldsymbol{Q}, t) = \rho_s(\boldsymbol{Q}) f(\boldsymbol{Q}, t). \quad [2]$$

$f(\boldsymbol{Q}, t)$ can be shown to obey the backward equation

$$\tau \frac{\partial f(\boldsymbol{Q}, t)}{\partial t} = \frac{1}{2} \sum_a \left[(\theta_a - |\boldsymbol{\theta}| Q_a) \frac{\partial}{\partial Q_a} f(\boldsymbol{Q}, t) \right] + \frac{1}{2} \sum_{a,b} \left[Q_a (\delta_{a,b} - Q_b) \frac{\partial^2}{\partial Q_a \partial Q_b} f(\boldsymbol{Q}, t) \right]. \quad [3]$$

By considering a coalescent dual process a probabilistic expansion can be derived from the backward equation (App. B). This probabilistic expansion forms the basis of the approximate count dynamics Eqs. 6,7. It also suggest an efficient method to sample from the transition density used in the simulation of the pathogen dynamics (Sec. E). A decomposition of f in Jacobi polynomials, which are the eigenfunctions of Eq. 3, allows us to derive an efficient method to analytically solve the Bayesian prediction and update steps (App. C). We use this alternative method to benchmark the quality of the inference achieved by the approximate counting scheme (Fig. 2).

B. Dirichlet mixture expansion of the belief distribution. Let us assume that the belief distribution at some time $B(\boldsymbol{Q}, t)$ is a mixture of Dirichlet distributions

$$B(\boldsymbol{Q}, t) = \sum_{\boldsymbol{m}} c_{\boldsymbol{m}}(t) \mathcal{D}(\boldsymbol{Q}, \boldsymbol{m} + \boldsymbol{\theta}), \quad [4]$$

with $\boldsymbol{m} = (m_1, m_2, \cdots)$ for integer m_i and mixture coefficients $c_{\boldsymbol{m}} \geq 0, \sum_{\boldsymbol{m}} c_{\boldsymbol{m}} = 1$. In particular this assumption holds true for the prior belief at $t = 0$ which is described by the single component $\boldsymbol{m} = 0$. Both the update and prediction step of the inference procedure can be reduced to a much simpler procedure in terms of the mixture coefficients as derived below. Using the mixture assumption $f(\boldsymbol{Q}, t) = B(\boldsymbol{Q}, t)/\rho_s(\boldsymbol{Q})$ is a mixture $f(\boldsymbol{Q}, t) = \sum_{\boldsymbol{m}} c_{\boldsymbol{m}}(t) f_{\boldsymbol{m}}(\boldsymbol{Q})$ of functions

$$f_{\boldsymbol{m}}(\boldsymbol{Q}) = (|\boldsymbol{\theta}|)_{|\boldsymbol{m}|} \prod_i \frac{Q_i^{m_i}}{(\theta_i)_{m_i}}, \quad [5]$$

where $(a)_m = a(a+1)\cdots(a+m-1)$. We can interpret the mixture coefficients as a probability distribution and define an average of a quantity x as $\langle x \rangle = \sum_{\boldsymbol{m}} c_{\boldsymbol{m}} x_{\boldsymbol{m}}$. Using this notation the expected frequency is given by

$$\hat{\boldsymbol{Q}} = \langle (\boldsymbol{m} + \boldsymbol{\theta})/(|\boldsymbol{m}| + |\boldsymbol{\theta}|) \rangle. \quad [6]$$

To perform the update step Eq. 3 we calculate

$$Q_a f_{\boldsymbol{m}}(\boldsymbol{Q}) = \frac{\theta_a + m_a}{|\boldsymbol{\theta}| + |\boldsymbol{m}|} f_{\boldsymbol{m} + \boldsymbol{e}_a}(\boldsymbol{Q}). \quad [7]$$

Here and in the following we define $\boldsymbol{e}_i$ as the unit basis vector with i-th entry 1 and all other entries 0. It follows that the new belief distribution is still a mixture of Dirichlet distributions with

$$c_{\boldsymbol{m}+\boldsymbol{e}_a}(t^+) = \frac{\theta_a + m_a}{|\boldsymbol{\theta}| + |\boldsymbol{m}|} c_{\boldsymbol{m}}(t^-)/Z, \quad [8]$$

where $Z = \langle (\theta_a + m_a)/(|\boldsymbol{\theta}| + |\boldsymbol{m}|) \rangle_{c_{\boldsymbol{m}}(t^-)}$ is a normalization constant. To gain intuition about Eq. 8 let us consider the static environment limit $\tau \to \infty$. Applying Eq. 8 sequentially starting from the prior belief shows that the belief distribution remains composed of a single component ($c_{\boldsymbol{m}} = 1$ for $\boldsymbol{m} = \boldsymbol{m}^\star(t)$, $c_{\boldsymbol{m}} = 0$ otherwise) in this limit, where $\boldsymbol{m}^\star(t)$ is the vector of counts of how often the different pathogens have been encountered. This is a classical result in Bayesian statistics (2). The fact that the belief distribution stays within the Dirichlet family when performing Bayesian updates starting from a Dirichlet prior is known as the conjugacy property of the Dirichlet prior with regards to the categorical likelihood function. The interpretation of $\boldsymbol{m}$ as the number of observations of a particular category motivates thinking about the θ_i as pseudocounts, as according to Eq. 4 they are added to the counts of different categories to pull the estimate of the frequencies closer to the prior expectations.

The prediction step Eq. 4 asks for an application of the backward operator $\mathcal{A}^\dagger$ to $f(\boldsymbol{Q},t)$. As the backward equation is linear it acts independently on the different components. Applying $\mathcal{A}^\dagger$ to $f_{\boldsymbol{m}}$ yields

$$\mathcal{A}^\dagger f_{\boldsymbol{m}}(\boldsymbol{Q}) = \frac{1}{2}(|\boldsymbol{m}|+|\boldsymbol{\theta}|-1)\left[\sum_{i=1}^{K} m_i f_{\boldsymbol{m}-\boldsymbol{e}_i}(\boldsymbol{Q}) - |\boldsymbol{m}| f_{\boldsymbol{m}}(q)\right]. \quad [9]$$

By considering $\mathcal{A}^\dagger$ in Eq. 9 as acting on $\boldsymbol{m}$ we obtain a dual death process description of the Wright-Fisher diffusion (3). Transitions from $\boldsymbol{m} \to \boldsymbol{m}-\boldsymbol{e}_i$ happen at a rate $\frac{1}{2}(|\boldsymbol{m}|+|\boldsymbol{\theta}|-1)m_i$. The belief distribution continues to be a mixture of Dirichlet distributions with mixture coefficients that change as

$$\tau\frac{\mathrm{d}c_{\boldsymbol{m}}}{\mathrm{d}t} = \frac{1}{2}\left(|\boldsymbol{m}|+|\boldsymbol{\theta}|\right)\left(\sum_i (m_i+1)c_{\boldsymbol{m}+\boldsymbol{e}_i}\right) - \frac{1}{2}\left(|\boldsymbol{m}|+|\boldsymbol{\theta}|-1\right)|\boldsymbol{m}|c_{\boldsymbol{m}}. \quad [10]$$

Let us define the mean count vector $\langle\boldsymbol{m}\rangle = \sum_{\boldsymbol{m}} c_{\boldsymbol{m}}\boldsymbol{m}$. Its dynamical equation is

$$\frac{\mathrm{d}\langle\boldsymbol{m}\rangle}{\mathrm{d}t} = \sum_{\boldsymbol{m}} \boldsymbol{m}\frac{\mathrm{d}c_{\boldsymbol{m}}}{\mathrm{d}t}. \quad [11]$$

Some algebra leads to the surprisingly simple yet exact equation

$$\tau\frac{\mathrm{d}\langle\boldsymbol{m}\rangle}{\mathrm{d}t} = -\frac{1}{2}\langle(|\boldsymbol{m}|+|\boldsymbol{\theta}|-1)\boldsymbol{m}\rangle. \quad [12]$$

The probabilistic expansion allows the derivation of an efficient approximate scheme for inference. For peaked mixture distributions $c_{\boldsymbol{m}}$ we can to a good approximation invert the order of calculating the expectation value and the product in Eq. 12, i.e.

$$\langle(|\boldsymbol{m}|+|\boldsymbol{\theta}|-1)\boldsymbol{m}\rangle \approx (|\langle\boldsymbol{m}\rangle|+|\boldsymbol{\theta}|-1)\langle\boldsymbol{m}\rangle. \quad [13]$$

For peaked mixture distribution the update equation for the mean counts upon encountering pathogen a is approximately

$$\langle\boldsymbol{m}(t^+)\rangle \approx \langle\boldsymbol{m}(t^-)\rangle + \boldsymbol{e}_a, \quad [14]$$

and the expected frequencies are approximately

$$\hat{\boldsymbol{Q}} \approx \langle\boldsymbol{m}+\boldsymbol{\theta}\rangle/(|\langle\boldsymbol{m}+\boldsymbol{\theta}\rangle|). \quad [15]$$

Dropping the explicit notation for the average and replacing $\boldsymbol{n} \coloneqq \boldsymbol{m}+\boldsymbol{\theta}$ we obtain the following equations for the total counts:

$$n_a(t^+) = n_a(t^-)+1 \quad \text{when } a \text{ is encountered} \quad [16]$$

$$\tau\frac{\mathrm{d}n_a}{\mathrm{d}t} = -\frac{1}{2}\left(|\boldsymbol{n}|-1\right)\left(n_a-\theta_a\right) \quad \text{for all } a. \quad [17]$$

C. Spectral expansion of the belief distribution. As in App. B we decompose the belief distribution but now we decompose $f(\boldsymbol{Q},t)$ into the eigenfunctions $g_{\boldsymbol{n}}(\boldsymbol{Q})$ of the backward operator. This approach leads to simpler decoupled equations for the prediction step of the Bayesian inference. To simplify notations we consider the case $K=2$ but the derivation generalizes to arbitrary K (3). Note that to describe the dynamics of the frequency of any particular pathogen, we can lump together the frequencies of the other $K-1$ pathogens for the model we consider. In particular, as the rates of immigration of different types are independent of composition of the population, the dynamics of the chosen pathogen does not depend on distribution of pathogens among the other types. The dynamics of the i-th type for arbitrary K can thus be mapped to $K=2$ by setting $\theta_1 \coloneqq \theta_i$ and $\theta_2 \coloneqq \sum_{j\neq i}\theta_j$.

The steady-state distribution specializes from the Dirichlet distribution to the Beta distribution for $K=2$

$$\rho_s(q) = q^{\theta_1-1}(1-q)^{\theta_2-1}/Z, \quad [18]$$

where $Z = B(\theta_1,\theta_2) = \Gamma(\theta_1)\Gamma(\theta_2)/\Gamma(\theta_1+\theta_2)$ is the Beta function. The eigenfunction $g_{\boldsymbol{n}}$ for an eigenvalue λ_n needs to fulfill the backward equation (Eq. 3), which leads to

$$\frac{1}{2}q(1-q)\frac{\mathrm{d}^2 g_n}{\mathrm{d}q^2}(q) + \frac{1}{2}(-\theta_2 q + \theta_1(1-q))\frac{\mathrm{d}g_n}{\mathrm{d}q}(q) = -\tau\lambda_n g_n(q), \quad [19]$$

where $q = Q_1$. Up to a rescaling this equation is a Jacobi differential equation. The solutions of this differential equation are the modified Jacobi Polynomials

$$g_n(q) = P_n^{(\theta_2-1,\theta_1-1)}(2q-1), \quad [20]$$

where $P_n^{(a,b)}(x)$ is the n-th Jacobi polynomial with eigenvalue

$$\lambda_n = \frac{1}{2\tau} n(n+\theta_1+\theta_2-1). \tag{21}$$

Let us first state a number of properties of these polynomials which we will need later. The polynomials form an orthogonal system with respect to the weight function $q^{\theta_1-1}(1-q)^{\theta_2-1}$, i.e.

$$\int_0^1 \mathrm{d}q\, g_n(q) g_m(q) q^{\theta_1-1}(1-q)^{\theta_2-1} = \delta_{n,m}\Delta_n(\theta_1,\theta_2), \tag{22}$$

where the normalization coefficients $\Delta_n(\theta_1,\theta_2)$ are given by

$$\Delta_n(\theta_1,\theta_2) = \frac{\Gamma(n+\theta_1)\Gamma(n+\theta_2)}{(2n+\theta_1+\theta_2-1)\Gamma(n+\theta_1+\theta_2-1)\Gamma(n+1)}. \tag{23}$$

For $n \geq 1$ the polynomials are related by the recursion formula (4)

$$q g_n(q) = \phi_n^- g_{n-1}(q) + \phi_n^0 g_n(q) + \phi_n^+ g_{n+1}(q), \tag{24}$$

with coefficients

$$\phi_n^- = \frac{(n+\theta_1-1)(n+\theta_2-1)}{(2n+\theta_1+\theta_2-1)(2n+\theta_1+\theta_2-2)}, \tag{25}$$

$$\phi_n^0 = \frac{1}{2} - \frac{\theta_2^2-\theta_1^2-2(\theta_2-\theta_1)}{2(2n+\theta_1+\theta_2)(2n+\theta_1+\theta_2-2)}, \tag{26}$$

$$\phi_n^+ = \frac{(n+1)(n+\theta_1+\theta_2-1)}{(2n+\theta_1+\theta_2)(2n+\theta_1+\theta_2-1)}, \tag{27}$$

while for $n = 0$

$$q g_0(q) = \phi_0^0 g_0(q) + \phi_0^+ g_1(q), \tag{28}$$

with coefficients

$$\phi_0^0 = \frac{\theta_1}{\theta_1+\theta_2}, \tag{29}$$

$$\phi_0^+ = \frac{1}{\theta_1+\theta_2}. \tag{30}$$

Now let us write $f(q,t)$ as a linear combination of the eigenfunctions g_n of the backward equation as,

$$f(q,t) = 1 + \sum_{n=1}^{\infty} d_n(t) g_n(q), \tag{31}$$

with time-varying coefficients $d_n(t)$. The prediction step of the Bayesian inference leads through Eq. 19 to a simple exponential decay of the d_n,

$$\frac{\mathrm{d}d_n(t)}{\mathrm{d}t} = -\lambda_n. \tag{32}$$

The update step of the Bayesian inference leads through Eq. 24 to

$$d_n(t^+) = \chi_n^- d_{n-1}(t^-) + \chi_n^0 d_n(t^-) + \chi_n^+ d_{n+1}(t^-), \tag{33}$$

where $\chi_n^-, \chi_n^0, \chi_n^+$ are constants. These constants are normalized versions of the coefficients ϕ that appear in Eq. 24,

$$\chi_n^x = \begin{cases} \phi_n^x/(c_0^0 + d_1(t^-)c_1^-) & \text{if pathogen 1 encountered,} \\ (\delta_{x,0} - \phi_n^x)/(1 - c_0^0 - d_1(t^-)c_1^-) & \text{if pathogen 2 encountered,} \end{cases} \tag{34}$$

for x in $-, 0, +$, and where $\delta_{x,0} = 1$ for $x = 0$ and $\delta_{x,0} = 0$ otherwise. For efficient numerical computation note that the update step can be performed as a matrix multiplication of the triadiagonal matrix which has c_n^0 along the diagonal, c_{n+1}^- above the diagonal, and c_{n-1}^+ below the diagonal with the vector of coefficients $\boldsymbol{d}(t^-)$ followed by a normalization step. We finally note that the expected frequency is obtained from Eq. 31 as

$$\langle q(t) \rangle = c_0^0 + c_1^- d_1(t). \tag{35}$$

The timescale over which the pathogen frequencies change is set by the slowest timescale of the stochastic dynamics, i.e.

$$\tau_c = \frac{1}{\lambda_1} = \frac{2\tau}{\theta_1+\theta_2}. \tag{36}$$

This result generalizes to the general case $K > 2$ as

$$\tau_c = \frac{2\tau}{|\boldsymbol{\theta}|}. \tag{37}$$

D. Defining an effective timescale for learning in changing environments. The total number of counts $|\boldsymbol{n}|$ follows a piecewise deterministic decay process interspersed by updates at random times. We approximate this stochastic dynamics by the deterministic equation,

$$\frac{\mathrm{d}|\boldsymbol{n}|}{\mathrm{d}t} = -\frac{1}{2\tau}(|\boldsymbol{n}|-1)(|\boldsymbol{n}|-|\boldsymbol{\theta}|)+\lambda, \tag{38}$$

where we have replaced the stochastic jumps in the counts do to pathogen encounters by a source-term λ equal to the rate of such jumps. This differential equation is separable and can be solved, which for $|\boldsymbol{n}(0)| = |\boldsymbol{\theta}|$ yields

$$|\boldsymbol{n}(t)| = \frac{|\boldsymbol{\theta}|+1}{2} + \sqrt{\eta}\tanh\left(\frac{t\sqrt{\eta}}{2\tau} + \operatorname{artanh}\left(\frac{|\boldsymbol{\theta}|-1}{2\sqrt{\eta}}\right)\right), \tag{39}$$

for

$$\eta = 2\lambda\tau + \frac{(|\boldsymbol{\theta}|-1)^2}{4} \tag{40}$$

Considering the limit of long environmental correlation times, $2\lambda\tau \gg \frac{(|\boldsymbol{\theta}|-1)^2}{4}$, we obtain the simplified expression

$$|\boldsymbol{n}(t)| - |\boldsymbol{\theta}| \approx \sqrt{2\lambda\tau}\tanh\left(\frac{\lambda t}{\sqrt{2\lambda\tau}}\right). \tag{41}$$

If we further take the limit $2\lambda\tau \gg (\lambda t)^2$ of no attrition, then we recover the simple linear scaling of the expected number of encounters with time

$$|\boldsymbol{n}(t)| - |\boldsymbol{\theta}| \approx \lambda t. \tag{42}$$

The number of remembered encounters $|\boldsymbol{n}(t)| - |\boldsymbol{\theta}|$ relative to the effective number of present pathogens $|\boldsymbol{\theta}|$ controls how much memory improves protection (Fig. 3).

E. Efficient simulation by sampling from the transition density. The transition density of going from $\boldsymbol{x}$ to $\boldsymbol{y}$ in time t for the Wright-Fisher diffusion can be written as (3)

$$f(\boldsymbol{x},\boldsymbol{y},t) = \sum_{|\boldsymbol{l}|=0}^{\infty} q_{|\boldsymbol{l}|}^{|\boldsymbol{\theta}|}(t) \sum_{\{\boldsymbol{l}:|\boldsymbol{l}|fixed\}} \mathcal{M}(\boldsymbol{l},\boldsymbol{x})\mathcal{D}(\boldsymbol{y},\boldsymbol{\theta}+\boldsymbol{l}), \tag{43}$$

where $\mathcal{M}(\boldsymbol{l},\boldsymbol{x})$ is the multinomial distribution

$$\mathcal{M}(\boldsymbol{l},\boldsymbol{x}) = \binom{|\boldsymbol{l}|}{\boldsymbol{l}} x_1^{l_1}\cdots x_K^{l_K}, \tag{44}$$

and where $q_{|\boldsymbol{l}|}^{|\boldsymbol{\theta}|}(t)$ are the transition functions of a dual pure death process. This description of the transition density can be derived based on similar arguments to those made in App. B (3). The death process describes the loss of unmutated lineages going backward in time through coalescence and mutations. For small times $q_{|\boldsymbol{l}|}^{|\boldsymbol{\theta}|}(t)$ is asymptotically normal with mean $\mu(t)$ and variance $\sigma^2(t)$ (5, 6)

$$\mu(t) = \frac{2\eta}{t}, \tag{45}$$

$$\sigma^2(t) = \begin{cases} \frac{2\eta}{t}(\eta+\beta)^2\left(1+\frac{\eta}{\eta+\beta}-2\eta\right)\beta^{-2}, & \beta\neq 0, \\ \frac{2}{3t}, & \beta = 0, \end{cases} \tag{46}$$

$$\text{where } \beta = \frac{1}{2}(|\boldsymbol{\theta}|-1)t, \tag{47}$$

$$\text{and } \eta = \begin{cases} \frac{e^{\beta}}{e^{\beta}-1}, & \beta\neq 0, \\ 1, & \beta = 0. \end{cases} \tag{48}$$

To sample from the transition density function we thus proceed in three steps: Generate a normally distributed random number $|\boldsymbol{l}|$ according to the asymptotic distribution of $q_{|\boldsymbol{l}|}^{|\boldsymbol{\theta}|}(t)$. Then draw $\boldsymbol{l}$ from $\mathcal{M}(l,\boldsymbol{x})$. Finally, draw $\boldsymbol{y}$ from $\mathcal{D}(\boldsymbol{\theta}+\boldsymbol{l})$.

2. Induced repertoire dynamics

A. Dependence of fold change upon a pathogen encounter on sparsity. To understand how memory production depends on environmental sparsity we specialize Eq. 9 to the case of a uniform prior distribution. We then have $|\boldsymbol{n}| \approx K\theta$ and $Z^- \approx \sum_a 1/K^{1/(1+\alpha)} = K^{\alpha/(1+\alpha)}$, which for $K\theta \gg 1$ leads to $\kappa = 1/(K^{1+\alpha}\theta)$. The fold change upon an encounter of a pathogen starting from a naive repertoire $\tilde{P}_a^- = 1/K$ thus depends as follows on the sparsity of the environment,

$$\tilde{P}_a^+/\tilde{P}_a^- = [1+1/\theta]^{1/(1+\alpha)}. \tag{49}$$

B. Mean-field dynamics. Besides the large changes of the naive repertoire upon a primary infection there are situations in which the inferred distribution is changing in a more continuous manner, e.g. updating in the limit of many previous samples, or the prediction step. We thus now ask how small changes in the expected frequencies of pathogens $\hat{\boldsymbol{Q}}$ change the coverage $\tilde{\boldsymbol{P}}$. We assume that there is no cross-reactivity $f_{r,a} = \delta_{r,a}$, and consider power-law cost functions, where we have optimal receptor frequency distribution $P_r^\star = Q_r^\beta / Z$ with $\beta = 1/(1+\alpha)$. As a preliminary we calculate the Jacobian

$$\frac{\partial P_r^\star}{\partial \hat{Q}_{r'}} = \delta_{r,r'} \frac{1}{Z} \frac{\partial \hat{Q}_r^\beta}{\partial \hat{Q}_r} - \frac{Q_r^\beta}{Z^2} \frac{\partial Z}{\partial \hat{Q}_{r'}}, \tag{50}$$

$$= \beta P_r^\star \left[\frac{\delta_{r,r'}}{\hat{Q}_r} - \frac{P_{r'}}{\hat{Q}_{r'}} \right], \tag{51}$$

We can then show that the dynamics in terms of $P_r^\star$ follows,

$$\frac{\mathrm{d}P_r^\star}{\mathrm{d}t} = \sum_{r'} \frac{\partial P_r^\star}{\partial \hat{Q}_{r'}} \frac{\mathrm{d}\hat{Q}_{r'}}{\mathrm{d}t}, \tag{52}$$

$$= P_r^\star \left[\beta \frac{\mathrm{d} \ln \hat{Q}_r}{\mathrm{d}t} - \sum_{r'} P_{r'} \beta \frac{\mathrm{d} \ln \hat{Q}_{r'}}{\mathrm{d}t} \right], \tag{53}$$

which is of the form of a replicator equation

$$\frac{\mathrm{d}P_r^\star}{\mathrm{d}t} = P_r^\star \left[f_r - \bar{f} \right] \tag{54}$$

with "fitness" $f_r = \beta \frac{\mathrm{d} \ln \hat{Q}_r}{\mathrm{d}t}$ and mean fitness $\bar{f} = \sum_{r'} P_{r'} f_{r'}$.

Based on this general result we now analyze the dynamics of the repertoire due to the sequential Bayesian filtering. Equivalently to Eq. 14 the change of inferred distribution $\Delta\hat{\boldsymbol{Q}} = \hat{\boldsymbol{Q}}^+ - \hat{\boldsymbol{Q}}$ upon encountering antigen a is given by

$$\Delta\hat{\boldsymbol{Q}} = \frac{\boldsymbol{e}_a - \hat{\boldsymbol{Q}}}{|\boldsymbol{n}| + 1}, \tag{55}$$

where $\boldsymbol{e}_a$ is the unit vector with a-th entry one and all other zero. Asymptotically for large $|\boldsymbol{n}|$ every update has a small effect only, and we might consider a mean-field description. In this description we replace $\boldsymbol{e}_a$ by its expectation value $\boldsymbol{Q}$ and define an average rate of change per unit time by multiplying the update size by the frequency λ of pathogen encounters:

$$\frac{\mathrm{d}\hat{\boldsymbol{Q}}}{\mathrm{d}t} = \frac{\lambda}{|\boldsymbol{n}| + 1} (\boldsymbol{Q} - \hat{\boldsymbol{Q}}). \tag{56}$$

Here $\boldsymbol{Q}$ is the actual distribution of pathogens, and $\hat{\boldsymbol{Q}}$ are the expected frequencies of pathogens based on the immune system's internal belief. For the prediction step we have a dynamics for counts, which we can convert into a dynamics for the inferred distribution. We have $\hat{Q}_r = n_r/|\boldsymbol{n}|$ and a dynamics on counts given by Eq. 7. From there we obtain

$$\frac{\mathrm{d}\hat{Q}_r}{\mathrm{d}t} = \frac{1}{|\boldsymbol{n}|} \frac{\mathrm{d}n_r}{\mathrm{d}t} - \frac{n_r}{|\boldsymbol{n}|^2} \frac{\mathrm{d}|\boldsymbol{n}|}{\mathrm{d}t}, \tag{57}$$

$$= -\frac{(|\boldsymbol{n}| - 1)|\boldsymbol{\theta}|}{2\tau|\boldsymbol{n}|} (\hat{Q}_r - \hat{Q}_r^0), \tag{58}$$

where $\hat{Q}_r^0 = \theta_r/|\boldsymbol{\theta}|$ is the prior guess for the distribution. Taken together we have

$$\frac{\mathrm{d}\hat{\boldsymbol{Q}}}{\mathrm{d}t} = \gamma(t)(\boldsymbol{Q} - \hat{\boldsymbol{Q}}) - \delta(t)(\hat{\boldsymbol{Q}} - \hat{\boldsymbol{Q}}^0). \tag{59}$$

with the (time-varying) coefficients

$$\gamma(t) = \frac{\lambda}{|\boldsymbol{n}(t)| + 1}, \tag{60}$$

$$\delta(t) = \frac{(|\boldsymbol{n}(t)| - 1)|\boldsymbol{\theta}|}{2\tau|\boldsymbol{n}(t)|}. \tag{61}$$

The fitness in the replicator equation is then

$$f_r = \beta \left(\gamma(t) \frac{Q_r}{\hat{Q}_r} + \delta(t) \frac{\hat{Q}_r^0}{\hat{Q}_r} \right). \tag{62}$$

The fixed point of the dynamics in a static environment $\delta(t) = 0$ is the optimal repertoire as expected from the asymptotic optimality of Bayesian inference. Replacing $\hat{Q}_r = (ZP_r^\star)^{1+\alpha}$ we then obtain a fitness

$$f_r = \frac{\gamma(t)}{Z^{1+\alpha}} \frac{Q_r}{(P_r^\star)^{1+\alpha}} \quad [63]$$

which except for the prefactor is equivalent to the population dynamics proposed previously in (7). That work did not consider the prefactor that leads to a slowing down of the dynamics with time to reflect a tradeoff between new evidence and past experience. The prediction steps relaxes the inferred distribution towards the prior distribution with a speed that for large $|\boldsymbol{n}|$ is proportional to $|\boldsymbol{\theta}|/\tau$.

C. Updating a cross-reactive repertoire. We now consider the repertoire dynamics in the presence of cross-reactivity. In a first order Taylor expansion the change in the repertoire composition upon a pathogen encounter is given by

$$\Delta \boldsymbol{P}^\star \approx \boldsymbol{J_G}(\hat{\boldsymbol{Q}}) \Delta \hat{\boldsymbol{Q}}, \quad [64]$$

where

$$(\boldsymbol{J_G})_{r,r'}(\hat{\boldsymbol{Q}}) = \frac{\partial G_r(\hat{\boldsymbol{Q}})}{\partial Q'_r} \quad [65]$$

is the Jacobian of the mapping function (Eq. 1).

The mapping between pathogen frequencies and the optimal repertoire takes the form $\boldsymbol{G}(\boldsymbol{Q}) = \boldsymbol{F}^{-1}\tilde{\boldsymbol{P}}^\star(\boldsymbol{Q})$ (if achievable given the constraint that no receptor frequency can be negative). $\tilde{\boldsymbol{P}}^\star(\boldsymbol{Q})$ is a function that depends on the cost function. The Jacobian can thus be calculated using the chain rule as

$$\boldsymbol{J_G}(\hat{\boldsymbol{Q}}) = \boldsymbol{F}^{-1}\boldsymbol{J}_{\tilde{\boldsymbol{P}}^\star}(\hat{\boldsymbol{Q}}). \quad [66]$$

For the power-law cost function we have $\tilde{P}_a^\star = R\hat{Q}_a^\beta/Z$, where $R = \sum_a f_{r,a}$ is the row sum of F, which assume to be constant. Analogously to the derivation of the Jacobian in the previous section we derive

$$\frac{\partial \tilde{P}_a^\star}{\partial \hat{Q}_{a'}} = \beta \tilde{P}_a^\star \left[\frac{\delta_{a,a'}}{\hat{Q}_a} - \frac{\tilde{P}_{a'}}{R\hat{Q}_{a'}} \right], \quad [67]$$

from which with some algebra follows

$$\Delta \boldsymbol{P}^\star \approx \left(\frac{R}{\tilde{P}_a^\star} \right)^\alpha \frac{1}{(1+\alpha)(|\boldsymbol{n}|+1)Z^{1+\alpha}} \left(R\boldsymbol{F}^{-1}\boldsymbol{e}_a - \boldsymbol{P}^\star \right). \quad [68]$$

Here, there is a departure from the dynamics of the number N_r of lymphocytes with receptor r proposed in (7),

$$\Delta N_r = N_r \Delta t \left[A(\tilde{N}_a) f_{r,a} - d \right], \quad [69]$$

where proliferation is proportional to $f_{r,a}$, instead of $(F^{-1})_{r,a}$.

We can also consider autonomous dynamics for the total number of lymphocytes carrying receptor r, $\boldsymbol{N} = \chi \boldsymbol{F}^{-1}\boldsymbol{n}^\beta$, with $\boldsymbol{P} = \boldsymbol{N}/|\boldsymbol{N}|$ (with no concern for the normalization of $\boldsymbol{N}$ to be constant). The population dynamics then read:

$$\Delta \boldsymbol{N} = B\boldsymbol{F}^{-1}\left[(\boldsymbol{FN})^{-\alpha}.\boldsymbol{e}_r \right] \quad \text{when } r \text{ is encountered} \quad [70]$$

$$\tau_N \frac{d\boldsymbol{N}}{dt} = \chi^{1+\alpha}\boldsymbol{F}^{-1}\left[(\boldsymbol{FN})^{-\alpha}.\boldsymbol{\theta} \right] - \boldsymbol{N}. \quad [71]$$

with $B = \chi^{1+\alpha}/(1+\alpha)$ and $\tau_N = 2(1+\alpha)\tau/(|(\boldsymbol{FN}/\chi)^{-\alpha}|-1)$ is a global homeostasis parameter. Powers are taken element-wise, and $\boldsymbol{x}.\boldsymbol{y}$ denotes the element-wise product of vectors. The prediction step Eq. (71) can be plausibly implemented by a thymic source defined as $\tau_N^{-1}\chi^{1+\alpha}\boldsymbol{F}^{-1}\left[(\boldsymbol{FN})^{-\alpha}.\boldsymbol{\theta} \right]$, and a constant decay term. The update step Eq. (70) could be realized by a quorum sensing mechanisms (through cytokine-mediated cell-cell communication) combined with differentiated responses as a function of affinity $f_{a,r}$.

3. Inference of high-dimensional categorical distributions from few samples

A. Mean cost versus time. In this Appendix we will derive analytical expressions for the optimized cost as a function of time

$$c(t) = \sum_{a=1}^{K} Q_a(t) c(\tilde{P}_a^\star(t)) \quad [72]$$

with the following simplifying assumptions: absence of cross-reactivity, $f_{a,r} = \delta_{a,r}$ and $\tilde{P}_a^\star(t) = P_a^\star(t)$; no attrition, $\tau \to +\infty$; and a power-law cost function $c(P) = P^{-\alpha}$. The prior on $\boldsymbol{Q}$ is a homogeneous Dirichlet distribution:

$$P(\boldsymbol{Q}) \propto \prod_{a=1}^{K} Q_a^{\theta-1}. \quad [73]$$

This problem is equivalent to the Bayesian inference of a distribution drawn from a Dirichlet meta-distribution. Asymptotic convergence properties of Bayesian inference procedures are well-established (2), but the convergence speed of Bayesian estimators of the distribution in the non-asymptotic regime has been much less studied to our knowledge. Analysing the behaviour of $c(t)$ is equivalent to analysing the convergence of the estimated distribution to the true one with increasing number of samples. Here we will establish the relevant scaling for few samples.

We consider the biologically relevant regime of high dimension but effective sparsity of the distribution, $K\theta \gg 1$, $\theta \ll 1$. Our main insight is that for such sparse distribution Bayesian inference is effective when the number of samples is on the order of a few $K\theta$, instead of the potentially much larger K.

The prominent role sparsity plays in allowing for more efficient estimation is reminiscent of compressed sensing (8). Non-asymptotic results about inference in high-dimensional settings have been explored recently in the context of machine learning (9). Both connections merit further exploration.

We define the expected cost as $\langle c(t)\rangle$, where the average is taken over both random choices of $\boldsymbol{Q}$, and random realizations of the pathogen encounters, $\boldsymbol{n}$, which are distributed according to:

$$n_a(t) \sim \text{Poisson}(\lambda t Q_a), \tag{74}$$

where λ is the encounter rate. The number of encounters determine the average belief for $\boldsymbol{Q}$,

$$\hat{Q}_a(t) = \frac{\theta + n_a}{K\theta + \sum_{a=1}^{L} n_a} \approx \frac{\theta + n_a}{K\theta + \lambda t}, \tag{75}$$

which itself shapes the optimal response and thus the cost through Eq. 72:

$$P_a^\star(t) = \frac{\hat{Q}_a(t)^{1/(1+\alpha)}}{Z}, \tag{76}$$

where Z is a normalization constant.

Note that the cost can be expressed in terms of a divergence between the best receptor distribution $\boldsymbol{P}^\star$ given full knowledge of $\boldsymbol{Q}$ and the actual receptor distribution $\boldsymbol{P}$. Defining $P_a^\star = Q_a^{1/(1+\alpha)}/\tilde{Z}$ and replacing into Eq. 72, we obtain:

$$c(t) = \tilde{Z}^{1+\alpha} \sum_a (P_a^\star)^{1+\alpha} P_a^{-\alpha} = c_\infty \exp\left[\alpha D_{1+\alpha}(\boldsymbol{P}^\star \| \boldsymbol{P})\right], \tag{77}$$

where $c_\infty = \tilde{Z}^{1+\alpha}$ is the asymptotic cost for $P = P^\star$ and where $D_\beta(\boldsymbol{P}\|\boldsymbol{Q}) := (\beta-1)^{-1} \ln[\sum_a P_a^\beta Q_a^{1-\beta}]$ is the Rényi divergence of order β, which reduces to the standard Kullback-Leibler divergence for $\beta = 1$, i.e. $\alpha = 0$.

B. Reducing the problem to a single pathogen. By symmetry all terms in the sum of Eq. 72 are equal on average and we have

$$\langle c(t)\rangle = K \int dq\, \rho(q) q \, \langle c(p)\rangle_n, \tag{78}$$

where we have introduced the short hand notations $q \coloneqq Q_a$ and $p \coloneqq \tilde{P}_a$, and where the average is taken over $n \coloneqq n_a$. The pathogen frequency q is approximately Gamma-distributed:

$$\rho(q) \approx \frac{(K\theta)^\theta}{\Gamma(\theta)} q^{\theta-1} e^{-K\theta q}. \tag{79}$$

In general, the expectation value depends through p on all previous encounters with any of the pathogens (i.e. all the other $n_{a'}$, $a' \neq a$). In high dimensions we can approximate this dependence by neglecting the correlation of the normalization factor Z with $\hat{q}$ and using an effective Z. For the power law cost functions we then have

$$p = \hat{q}^{\frac{1}{1+\alpha}}/Z \quad \text{with } Z \approx K\langle \hat{q}^{\frac{1}{1+\alpha}}\rangle. \tag{80}$$

For logarithmic cost this simplifies to $p = \hat{q}$ and $Z = 1$. From Eq. 80 it follows that

$$\langle c(t)\rangle \approx K Z^\alpha \langle q \hat{q}^{-\frac{\alpha}{1+\alpha}}\rangle = K^{1+\alpha} \langle \hat{q}^{\frac{1}{1+\alpha}}\rangle^\alpha \langle q \hat{q}^{-\frac{\alpha}{1+\alpha}}\rangle. \tag{81}$$

We have $\hat{q} = (\theta + n)/(K\theta + \lambda t)$, but as the equation is invariant to a linear rescaling of $\hat{q}$ we can replace $\hat{q}$ by simply $\theta + n$.

C. Costs for perfect or no information. Asymptotically the distribution is learned perfectly and we have $\hat{q} = q$. Plugging this into the expressions derived previously we obtain

$$\bar{c}_\infty := \langle c(\infty)\rangle \approx K^{1+\alpha}\langle q^{1/(1+\alpha)}\rangle^{1+\alpha}, \quad \bar{c}_\infty \approx -K\langle q\ln(q)\rangle, \tag{82}$$

for the power-law and logarithmic cost function respectively. Performing the integrals we obtain

$$\bar{c}_\infty \approx (K\theta)^\alpha \Gamma(1/(1+\alpha))^{1+\alpha}, \quad \bar{c}_\infty \approx \ln(K\theta) + \gamma, \tag{83}$$

where $\Gamma(z)$ is the Gamma function and γ the Euler-Mascheroni constant. For $\alpha = 1$ this specializes to $\bar{c}_\infty = \pi K\theta$. For the logarithmic cost $\bar{c}_\infty$ is equal to the Shannon entropy of the distribution, which suggests an interpretation of $K\theta$ as the effective number of pathogens that are present.

We can compare these costs for those obtained for a uniform repertoire

$$\bar{c}_0 := \langle c(0)\rangle = K^\alpha, \quad \bar{c}_0 = \ln(K), \tag{84}$$

to obtain

$$\bar{c}_\infty/\bar{c}_0 \propto \theta^\alpha, \quad \bar{c}_\infty/\bar{c}_0 = (\ln(K\theta) + \gamma)/\ln(K) \tag{85}$$

for power law and logarithmic cost respectively. As expected, in more sparse environments a larger relative improvement can be obtained by learning the distribution.

D. Scaling in the limit of few samples. In the limit of small sampling, each pathogen has been seen at most once, meaning that n is binary and distributed according to a Bernoulli variable with mean $\lambda t q$. Then we can use the approximation $\langle(\theta + n)^\beta\rangle \approx \theta^\beta(1 - \lambda t q) + \lambda t q$ to obtain

$$\langle \hat{q}^{\frac{1}{1+\alpha}}\rangle = \theta^{\frac{1}{1+\alpha}} + (1 - \theta^{\frac{1}{1+\alpha}})\lambda t\langle q\rangle \approx \theta^{\frac{1}{1+\alpha}} + \frac{\lambda t}{K} \tag{86}$$

$$\langle q\hat{q}^{-\frac{\alpha}{1+\alpha}}\rangle = \theta^{-\frac{\alpha}{1+\alpha}}\langle q\rangle + (1 - \theta^{-\alpha}1 + \alpha)\lambda t\langle q^2\rangle \approx \frac{\theta^{-\frac{\alpha}{1+\alpha}}}{K}\left[1 - \frac{\lambda t}{K\theta}\right] \tag{87}$$

Putting things together we obtain

$$\langle c(t)\rangle \approx K^\alpha\left(1 + \frac{\lambda t}{K\theta}\theta^{\frac{\alpha}{1+\alpha}}\right)^\alpha\left(1 - \frac{\lambda t}{K\theta}\right) \approx K^\alpha\left(1 - \left(1 - \alpha\theta^{\frac{\alpha}{1+\alpha}}\right)\frac{\lambda t}{K\theta}\right), \tag{88}$$

which, except for a correction that vanishes as $\theta \to 0$, scales with $\lambda t/(K\theta)$.

For the logarithmic cost we approximate similarly $\langle\ln(\theta + n)\rangle \approx (1 - \lambda t q)\ln(\theta) + \lambda t q\ln(1 + \theta) \approx (1 - \lambda t q)\ln(\theta)$. We then have

$$\langle c(t)\rangle \approx \ln(K\theta + \lambda t) - K\ln(\theta)\left[\langle q\rangle - \lambda t\langle q^2\rangle\right]. \tag{89}$$

Using the formulas for the first and second moments we obtain

$$\langle c(t)\rangle \approx \ln(K\theta + \lambda t) - \ln(\theta) - \frac{\lambda t}{K\theta}\ln(1/\theta) \tag{90}$$

Approximating further we have

$$\langle c(t)\rangle \approx \ln(K)\left(1 - \left(1 - \frac{\ln(K\theta) - 1}{\ln(K)}\right)\frac{\lambda t}{K\theta}\right). \tag{91}$$

Again the relative cost depends solely on $\lambda t/(K\theta)$ except for logarithmic corrections that vanish as $K \to \infty$ for fixed $K\theta$.

4. Infection cost in the expansion-delay regime

We have previously described mechanistic models that give rise to a power-law dependency of infection cost on the coverage (7), in which we have assumed thatthe crucial determinant of infection cost is set by the time delay to recognition of the pathogen by the immune system. Experimental evidence shows that the initial recruitment of a large fraction of all specific lymphocytes often happens rapidly compared to the time it takes for the adaptive immune system to start clearing the infection (10). We thus might hypothesize that the advantage of higher precursor numbers lies not in shortening the time to detection but in shortening the time to response by a sufficiently large number of effector cells.

To derive the scaling of infection cost with coverage under these conditions, we consider that after an infection at time 0, the number of specific cells grows exponentially with a rate γ, $N(t) = N(0)e^{-\gamma t}$. During the same time the pathogen population grows exponentially as well at a rate γ_p, $P(t) = P(0)e^{-\gamma_p t}$ until a time $t^\star$ at which a threshold level $N^\star$ of specific cells is reached. The expansion-delay time scales as $t^\star = \ln(N^\star/N(0))/\gamma$. If we assume that the cost of an infection is proportional $P(t^\star)$, then the cost scales as a power law with the initial number of specific cells $N(0)^{-\gamma_p/\gamma}$.

References

1. Etheridge A (2012) *Some Mathematical Models from Population Genetics.* (Springer-Verlag), Lecture no edition, p. 119.
2. Gelman A, Carlin JB, Stern HS, Rubin DB (2004) *Bayesian data analysis.* (CRC Press, Boca Raton).
3. Griffiths RC, Spano D (2010) Diffusion processes and coalescent trees. *arXiv Prepr. arXiv1003.4650.*
4. Song YS, Steinrücken M (2012) A simple method for finding explicit analytic transition densities of diffusion processes with general diploid selection. *Genetics* 190(3):1117–1129.
5. Griffiths RC (1984) Asymptotic line-of-descent distributions. pp. 67–75.
6. Jenkins PA, Spano D (2015) Exact simulation of the Wright-Fisher diffusion. *arXiv* 27(1):23.
7. Mayer A, Balasubramanian V, Mora T, Walczak AMA (2015) How a well-adapted immune system is organized. *Proc. Natl. Acad. Sci.* 112(19):5950–5955.
8. Zdeborová L, Krzakala F (2015) Statistical physics of inference: Thresholds and algorithms. *arXiv Prepr. arXiv1511.02476* 8732(i):1–62.
9. Advani M, Ganguli S (2016) Statistical mechanics of optimal convex inference in high dimensions. *Phys. Rev. X* 6(3):1–16.
10. Zehn D, Lee SY, Bevan MJ (2009) Complete but curtailed T-cell response to very low-affinity antigen. *Nature* 458(7235):211–214.

5. Supplementary figures

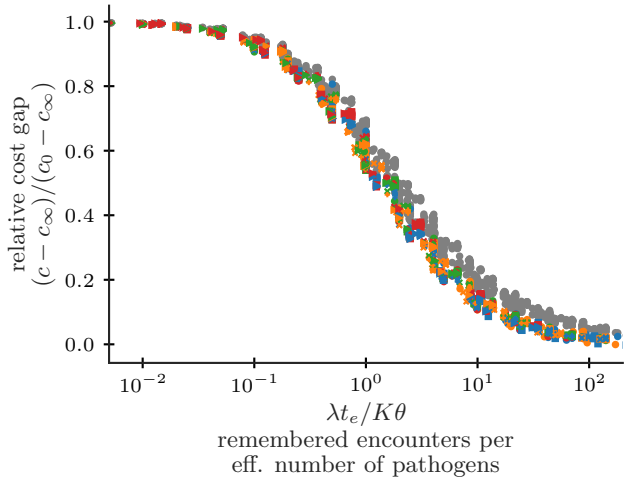

Fig. S1. Comparison of the dynamic scaling of the relative cost gap for $c(\tilde{P}_a) = -\ln \tilde{P}_a$ (colored) and $c(\tilde{P}_a) = 1/\tilde{P}_a$ (grey, see Fig. 3). For both cost functions rescaling parameters as $\lambda t_e / K\theta$ collapses the relative cost gaps for different parameters onto a similar master curve.

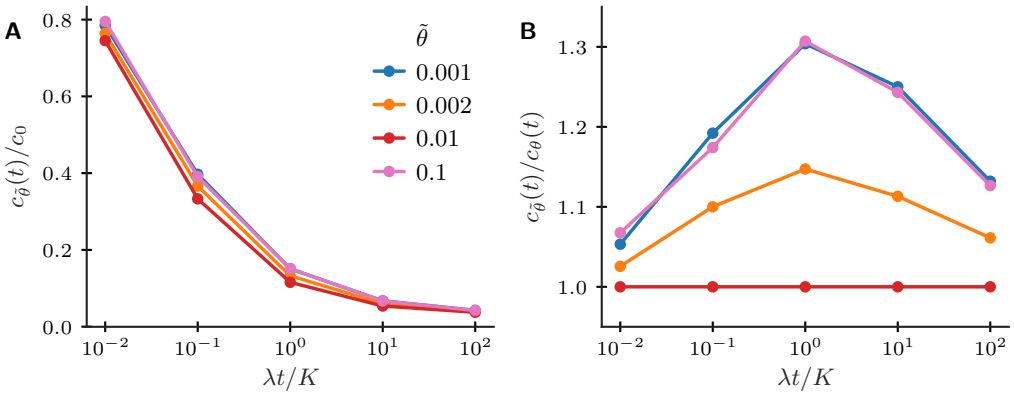

Fig. S2. Most of the benefit of memory is still achieved if the immune system has a slightly wrong prior about the antigen distribution sparsity θ. (A) Scaling of mean cost of infections with different priors $\tilde{\theta}$ for a correct $\theta = 0.01$ ($K = 1000$). (B) Relative increase in cost by using a wrong prior vs age.

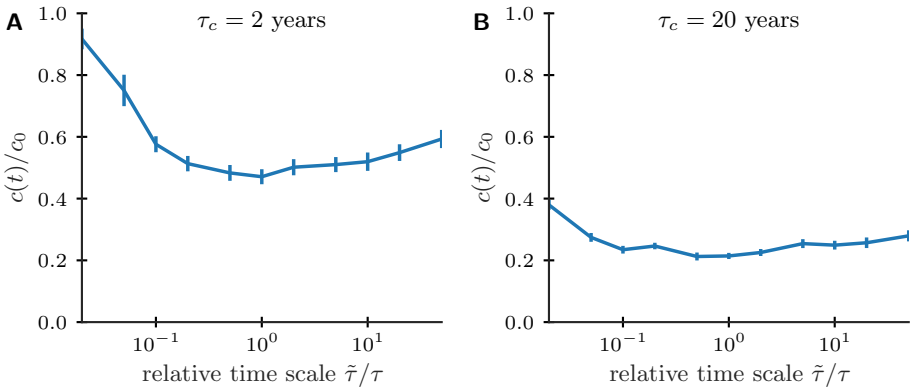

Fig. S3. The timescale of memory attrition does not need to be finely tuned for near-optimal prediction. Relative cost of infection at age 40 years as a function of the relative attrition timescale $\tilde{\tau}/\tau$ for environments with different correlation times $\tau_c = 2\tau/K\theta$. Parameters: $\alpha = 0.5$, K, θ, λ as in Fig. 5.

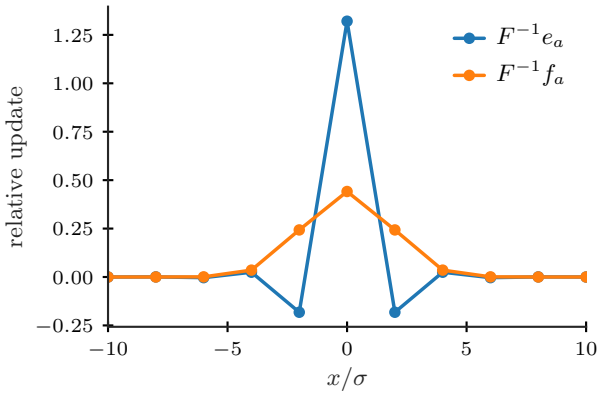

Fig. S4. While the optimal update of the receptor frequencies for uncorrelated pathogen environments (blue line) shows a depletion of close but suboptimal receptors such an effect might be reversed for correlated pathogen environments (orange line). We use a Gaussian kernel to describe cross-reactivity $f(|a-r|) = e^{|a-r|^2/(2\sigma^2}$, and a spacing of receptors 0.5σ. Assuming a change in inferred pathogen frequencies upon encountering pathogen a of e_a for the uncorrelated case (unit vector in a direction) and $f_a = e^{(x-a)/(2\sigma_c^2)}$ with $\sigma_c = 2\sigma$ for the correlated case we obtain the optimal update in terms of receptor frequencies by multiplying the inverse of the cross-reactivity matrix F with the change in frequencies.

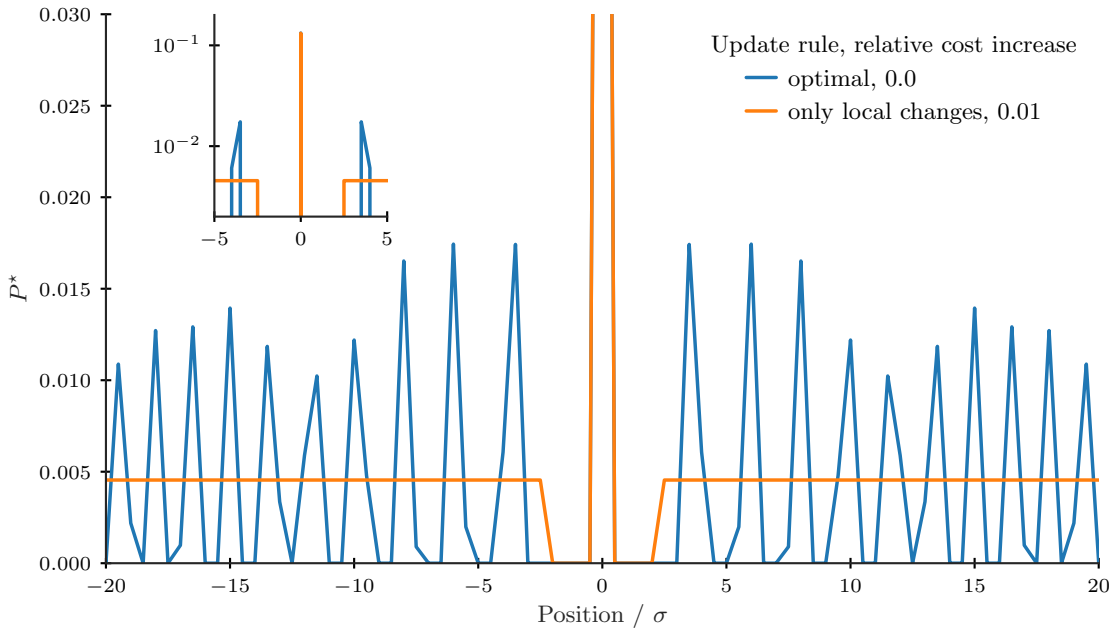

Fig. S5. For large changes in inferred frequencies the optimal repertoire can change drastically (blue line). A restriction to local changes yields a repertoire providing close to optimal protection, which exhibits a depletion of good, but suboptimal receptors (orange line). The inset shows a zoomed view of the distribution around the position of the sampled pathogen. To obtain the global (local) optimal repertoire we use a projected gradient algorithm to minimize the cost function over all receptor distributions (over only the disribution of receptors within 2σ of the pathogen in shape space). We use a logarithmic cost function, $|\boldsymbol{\theta}| = 10$, Gaussian kernel $f(|a-r|) = e^{|a-r|^2/(2\sigma^2)}$, and a spacing of receptors 0.5σ.